Development of ASHRAM; A new Human-Reliability-Analysis Method for Aviation Safety

Dwight P. Miller, Ph.D., CPE
Systems Reliability Department
Sandia National Laboratories*
Albuquerque, New Mexico

ABSTRACT

The primary purpose of the Aviation Safety Human Reliability Analysis Method, or ASHRAM is to predict plausible aviation-accident scenarios before they occur. An underlying premise of ASHRAM, is that many significant human errors can occur as a result of a combination of situational factors, or "error-forcing context" that can trigger cognitive 'error mechanisms' in personnel., which can lead to the execution of unsafe acts. The method allows aviation researchers to analyze accidents and incidents retrospectively, by answering questions and filling in forms, or prospectively, by systematically generating families of plausible scenarios based on a small set of initiators. ASHRAM is packaged in a brief, readable format, with step-by-step instructions, and with real-world examples so that it can be utilized by a variety of researchers, modelers, analysts, trainers, and pilots with a variety of backgrounds.

INTRODUCTION

Since the late 1950s, Sandia National Laboratories (SNL) has played a leadership role in the development of human reliability analysis (HRA) techniques for high-risk/consequence operations. The most recent of these is ASHRAM, which gets its roots from "A Technique for Human Event ANAlysis," or ATHEANA, which was developed for the US Nuclear Regulatory Commission [1] by SNL and other laboratories. Due to the field-proven usefulness of ATHEANA, and the Clinton administration's initiative to improve commercial airline safety tenfold, SNL spent internal research and development funds to develop ASHRAM during FY99 and FY00. The current form is ready for beta testing and refinement. ASHRAM is best applied by a team of diverse experts, but can also be used by a single analyst.

This paper summarizes parts, but not all of the ASHRAM project technical report [2]. First, the cognitive model will be described, followed by a summary of the procedures to perform retrospective and prospective analyses. Conclusions will address unique benefits to be derived from ASHRAM usage and appraises potential future directions for the technique.

UNDERLYING MODELS

ASHRAM utilizes a simplified three-stage cognitive model, which may be more intuitive for non-cognitive psychologists (see Fig. 1.) The model is not intended to describe behavior, but helps to categorize factors that influence the behavior of the operators. The three rounded boxes in the middle of the figure depict environmental perception (EP), which includes perceptual processes, attention, detection, recognition, monitoring, and overall understanding of the state of the aircraft/environment system); reasoning and decision-making (R/D/M), which includes thinking, judgments, remembering training, diagnosis, response selection, and creative problem-solving; and action (A), which includes control inputs to airframe, communications to crew, etc. The remainder of Fig. 1 shows the interrelationships of the environment (traffic and weather) and the aircraft condition (AC) with input channels to the pilot (displays, radio, communications, etc.), and how operator factors (OFs), such as stress and fatigue, and design factors (DFs), such as engine performance, can influence the three stages of processing.

The theoretical underpinnings of how unsafe actions (UAs) occur comes directly from ATHEANA. An UA is an overt action inappropriately taken or omitted that violates a critical flight function (CFF), such as thrust, attitude control, or airframe integrity. The term UA avoids any inference of blame and accommodates the notion that people are often "set up" by circumstances to make actions that are unsafe. This error-forcing context (EFC) is the combined effect of aircraft conditions (ACs), operator and design-based performance shaping factors (PSFs), procedural factors (PFs), weather (WX), traffic (TF), and CRM issues. In these circumstances, the crew does not knowingly commit and error, they perform "correct" actions, as they seemed to be at the time. Contributory actions (CAs) precipitate or lead to the

UA, but are not in and of themselves necessarily inappropriate or unsafe. Error mechanisms (EMs) are the cognitive processes that have been cultivated over time to deal with environmental demands that may tax limited processing resources, but when employed inappropriately or out of context. can contribute to UAs. An example is the expectancy bias. It helps us anticipate frequently experienced events and perform most efficiently when correct, but can interfere with our ability to deal with unexpected events. Error mechanisms mostly apply to the R/D/M stage of the cognitive model, however a few can apply to both the EP and A stages.
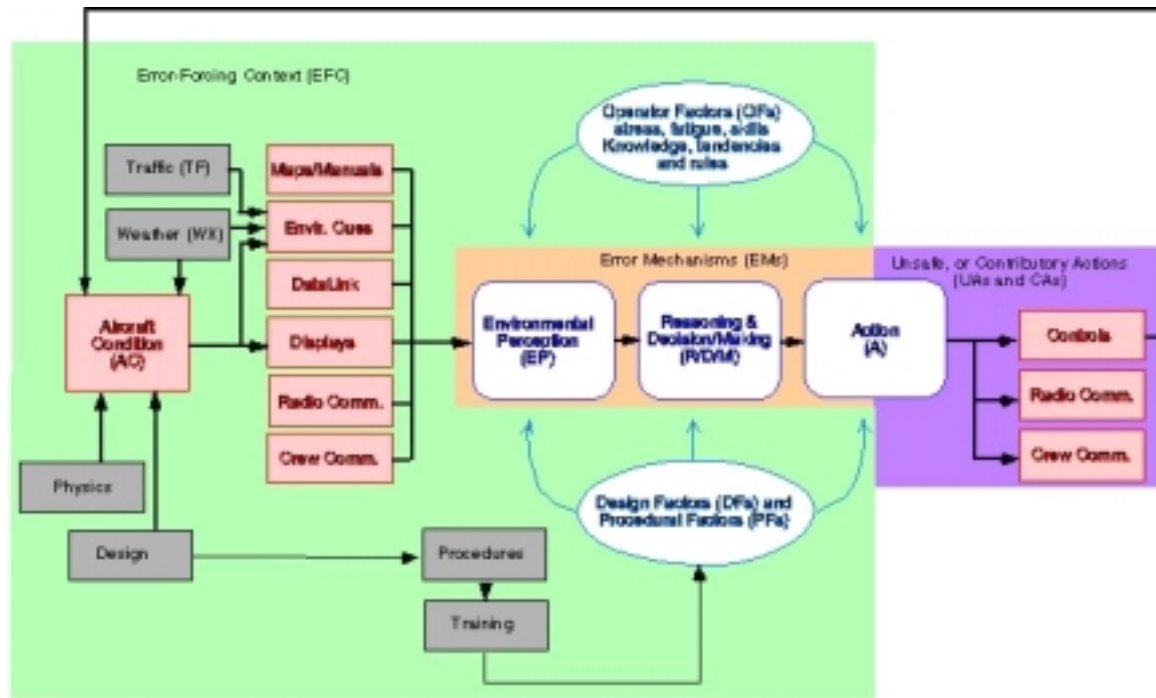

Figure 1. Human-system interaction model for ASHRAM

RETROSPECTIVE ANALYSIS

One purpose of the retrospective analysis is to unravel the event anew, using ASHRAM model constructs and terminology to ultimately uncover and describe the error-forcing context (EFC) that contributed to the event. This process will hopefully lead to an increased understanding of the ways in which the pilots (or ATC personnel) were lured into making inappropriate responses. Another purpose is to feed the prospective-analysis methodology by providing an issue for a base-case scenario, from which variations are spawned, representing the many ways things can 'go wrong.' The methodology supplies a series of questions and forms to be filled by the analyst/team, using available accident documentation such as incident reports, NTSB reports, etc. This activity can be accomplished by an individual or a team. Please refer to [2] for more information on the retrospective analysis.

PROSPECTIVE ANALYSIS

The heart of ASHRAM, and the bulk of the technical report [2] addresses the prospective-analysis methodology, where the goal is to predict plausible accident scenarios that have not yet occurred. More specifically, this process identifies elements of EFCs that contribute to unsafe actions, analyzes situations where pilots perform actions not required for emergency response, and documents families of related undesirable events. A prospective analysis should take a team from one week to several months to complete, depending on available materials, issue complexity, and teamwork. Recommended expertise

includes aviation safety, piloting, ASHRAM methodology; and optionally ATC operations, airframe technical systems, and weather. An overview of the process is depicted in Fig. 2, which can be used as guidance for the following section.
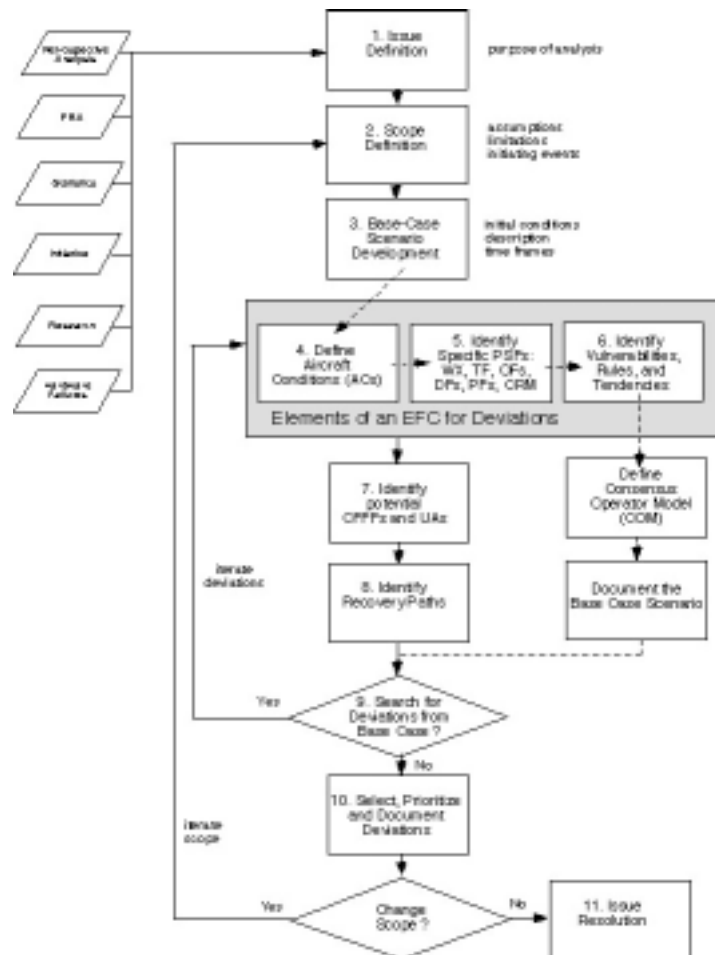

Figure 2. ASHRAM prospective-analysis flowchart

Process

Step 1. The first step is to define the core issue, or purpose of the analysis. The issue can come from a variety of sources, including a retrospective analysis, probabilistic risk assessments, government statistics, safety management programs, newspaper headlines, etc. The analyst who looks to specific accidents for inspiration, needs to step back from them to abstract a more generic class of events, based on similar initiating events, flight circumstances, and EFCs. This is a very important, necessary part of the issue-development process! An example of an issue would

be 'crew experiences with partial engine failures during flight.'

Step 2.. This step limits the scope of the analysis by setting additional boundaries of concern around the issue, including initiators, assumptions, system-related initial conditions, critical flight functional failures, and possible sets of human responses. The scope is best determined at the beginning of the prospective analysis, and should remain constant for the base-case scenario and the development of its plausible deviations (inner loop of Figure 2). Suggested parameters of scope include, but should not be limited to: type or series of aircraft; number,

and/or experience level of cockpit crewmembers; phase of flight, or specific operation, such as "land and hold short"; ground operations; CFFs compromised; nature or class of the initiating event; cockpit workload level; simultaneity of events; etc. As deviation scenarios are teased out and exhausted, the scope can be changed, and then iteratively revised (outer loop of Fig. 2). An important element in defining the scope of the prospective analysis is establishing initiating events. In addition to providing domain boundaries, initiating events can also spark the generation of base-case or deviation scenarios. Typically, in aircraft accidents, an initiating event is a mechanical, electrical, or chemical failure or change of state that requires responses on the part of the crew. Human actions can also be initiating events. Another potentially helpful way of considering initiating events, is by looking at classes of initiating events and examples within each class. For example, a broken fan blade may be an example of the class internal engine failures. This organization technique can help the analyst tease out an exhaustive list of potential ways that a particular CFF may be compromised.

Step 3. The base-case scenario represents the most realistic description of expected aircraft and crew behavior, and typically has a successful outcome. It provides a basis from which numerous deviation scenarios can be identified and described, hereafter referred to as simply deviations. It is the deviations that usually include UAs and can result in unsuccessful outcomes. Additional characteristics of base-case scenarios are: operationally well-defined, well-defined physics, may be well documented in public or proprietary references. Because the base-case scenario is based on a 'textbook case' and has a successful outcome, the progression through the flowchart in Fig. 2 differs from the progression that deviations follow. Steps 7 and 8 are bypassed, as they do not apply. Usually, one base-case scenario will be used in a prospective analysis at a time.

Because events can happen quickly in flying operations, it is advantageous to have a set of anticipated time frames. Each time frame is an estimated window of opportunity for a certain event to take place. Knowing which segments of a base-case scenario can happen quickly and which can happen more slowly, can help the analyst anticipate the information-processing constraints put on the crew. The analyst must also consider what the most relevant nominal aircraft conditions (ACs) will be over the time frames of interest. These might include: flight attitudes affected , status of certain instruments,

location of the crew, etc. Although there may be no definitive source of information for the ACs of interest, the appropriate experts on the team will have to make their best estimates of the behavior of relevant parameters over the time frames of interest.

At this point in the process, there is a divergence in procedure between base-case scenarios and deviations. For the base case, the EFC may be present, but the pilots respond correctly despite the factors that might otherwise entice them to perform UAs. The arrows indicating flow for the base case are dashed and "skip over" the shaded box and proceed to two boxes on the right-hand side to define the COM and document the base case. This loop is followed only once for each base-case and set of related deviations.

The most important component of the base-case scenario is the consensus operator model (COM). If a scenario is well defined and consistently understood among many pilots, the COM is the consensus, most-appropriate set of crew responses, and may be reflected in airline-published checklists. If actual, best practices deviate from published checklists, the best practices would prevail. The initiating event and the COM together form the basis for a base-case scenario.

Documentation for the base-case scenario should include: a description of initial conditions of the plane, flight, and crew; a list of assumed causes of the initiating event; a list of any other assumptions that are pertinent to the scenario; a brief, general description of the expected sequence of events, starting slightly before the initiating event; a description of the expected sequence and timing of aircraft behavior and responses; the expected trajectories of key flight parameters, plotted over time; and key pilot actions expected during the scenario progression.

Searching for deviant scenarios begins after base-case documentation with Step 9 and cycles up through Steps 4, 5, 6, 7, and 8 in an iterative manner until all deviant scenarios are exhausted for the given set of initial conditions, assumptions and limitations.

Step 4. After the decision to search for deviations has been made (in Step 9), new and different ACs need to be defined that change the situation and can potentially contribute to an EFC. Perhaps the most important source of variation to the EFC, the ACs need to be redefined as iterations of deviations proceed. As the way the plane responds to changes

from the base-case scenario script, the EFC changes, generating additional potential deviations and associated UAs. For example, an engine that had significant vibrations smoothes out and runs normally for the remainder of the flight. This change of AC may lead to the assumption that the engine is in good condition, when it, in fact, has problems.

Step 5. At this point, PSFs need to be considered as conditions that may make information-processing or action errors more likely. The PSFs, when combined with the ACs form the EFCs. The PSFs could be based on weather (WX), traffic (TF), operator factors (OFs), design factors (DFs), procedural factors (PFs), and crew-resource-management (CRM) issues. As PSFs and EFCs get identified, they will serve as fertile ground for UAs and CFFs discussed in Step 6.

Step 6. Not every pilot can be expected to know everything about his aircraft, its systems, their interrelationships, and all symptoms of all possible problems. This step attempts to identify any relevant gaps in the knowledge base associated with the base-case scenario, the behavior of its systems, relationships among its interacting systems, etc.

In addition, the aviation system uses hundreds of rules to keep it safe. Tendencies in pilots' behaviors are the most likely courses of action based on experience, knowledge, and rules. A tendency to respond in a particular manner to a situation may at first seem to be the most natural, comfortable set of decisions and actions. In some cases, the familiar response is so comfortable and automatic, that little R/D/M takes place. Tendencies are analogous to error mechanisms in that they are correctly applied most of the time, and save cognitive processing effort. However, when they are used in the wrong situation, they are considered errors in the R/D/M stage of information processing, and UAs can result.

The result of completing Steps 4 through 6 is a thorough description of the EFC for the scenario(s) being considered as deviations from the base case. Recall that a philosophical premise of the ASHRAM approach, is that significant human errors occur as a result of a combination of aircraft, airspace, weather, and crew conditions and other factors that trigger error mechanisms in the pilots. As the term EFC suggests, pilots can be tricked into executing UAs..

Step 7. The CFFFs, or critical flight function failures, are failures in the critically needed functions for safe flight. Any number of UAs can lead to a

CFFF. For example, the loss of thrust can be brought about by a number of UAs, including: throttling back power to engine, pulling fire extinguisher handle, turning engine fuel pump off.

Unsafe actions can be defined totally by context, where under one set of circumstances the action is not unsafe, but in another it becomes unsafe. An example is when a pilot changes altitude or heading and compromises airspace separation. Another source can be written procedures that are not 100% correct for all circumstances of use. Another source is taking instructions from another, as when a PIC asks the FO to perform some cockpit action. Three UA source paths can apply here: 1) the instruction is complied with, but is an unsafe action, 2) omitting the action, when it is the correct thing to do, and 3) performing the action incorrectly or incompletely. Although many HRA techniques differentially analyze behavior based on the type of error, as in errors of omission and errors of commission, ASHRAM concentrates on the context and the error mechanisms involved.

ASHRAM is flexible in its process in that it allows the analyst to follow two general paths--either generate UAs from variations in the EFC, or study the circumstances that may precipitate a given, pre-defined UA. The former approach is called a 'forward search' because it follows the flowchart in Figure 2 and the theoretical logic of UAs resulting from an EFC. Unlike many HRA techniques that need the unsafe actions as input to the method, and then identify relevant PSFs, or calculate human error probabilities given known scenarios, ASHRAM is somewhat unique in not only allowing for the generation of scenarios, but also the 'discovery' of UAs. The technique allows for the 'organic' germination of UAs directly from the initial conditions, initiating event, and the EFC. The approach takes the pressure off of the team to generate all possible outcomes of UAs earlier in the process. Forsythe and Wenner [3] have extolled the virtues of this "organic approach" to HRA. They see problems with generating every possible way that operators can make errors, and see advantages in enumerating the system conditions and characteristics that breed human errors.

If the CFFFs and UAs of interest are already defined, then the analyst documents them explicitly in this step, and performs a 'reverse search.' This search process consists of finding ACs, PSFs, and knowledge vulnerabilities, rules and tendencies that relate to and precipitate its manifestation. In a sense, the search is for elements of the EFC that set the stage

for pilot(s) to perform UAs. This is done by moving through the cognitive model backwards, by beginning in the A box and moving left to the R/D/M and EP boxes in a search to find what elements within the EFC could lead to error mechanisms that affect perception and reasoning.

Step 8. The cyclic process of generating and documenting a plausible deviation scenario needs to include enumerating plausible recovery paths that prevent the scenario from ending in a terminal event (one that signifies the unsuccessful termination of a flight). Recovery paths are limited to activities that take place after an UA has been committed. The overall likelihood of a deviation scenario proceeding toward a terminal-event conclusion is based on the probability of the unsafe act being committed underline(combined with) the probability that recovery does not occur. Finding ways to recover from UAs can be as involved a process as finding ways that UAs can occur.

Step 9. This is a decision node in the bottom half of Figure 2, which calls for the analyst to decide if another deviation is to be searched for by reverting back to Steps 4, 5, and 6 in order to make changes in the EFC. A deviation is a minor variation in some aspects that is otherwise based upon the base-case scenario. Although a different scenario altogether, complete with UAs, a deviation remains a 'family member' to a base-case scenario.

Step 10. As the cyclic process of iteration from Step 9 up and through Steps 4 - 8 continues, numerous deviant scenarios are accumulated. In order for them to be successful in passing through to documentation, they need to pass some fairly simple, straightforward criteria. The criteria can be any set deemed relevant by the team, but several will be suggested here: relevance to the issue, matching scope limitations, related to base case, uniqueness, plausibility of physics and human behavior. If any scenario does not meet the above criteria, it should be dropped from further consideration. It might well serve as a seed for another set of related scenarios, with a minor change in scope. If this is the case, after documentation, the analyst/team should consider iterating the scope.

For any given deviant scenario, two alternative means of documentation are suggested. First, is writing out the event in narrative format. This approach has the obvious advantages of including as much detail as desired and reading like a story. Unfortunately, if many deviations are forthcoming

from a prospective analysis, the writing can get laborious. An alternative is the event-tree style flow chart, where several possible deviant scenarios are outlined in a diagram showing their relationships that are based on decisions made or action taken. Its compact efficiency makes it desirable for families of deviations that are all minor variants of each other.

Change Scope? Having completed all of the possible deviations of the base-case scenario, there may be a desire to generate additional scenarios based on a shift in scope. If the issue remains consistent, the scope can be altered, a new base-case scenario can be written up, and new families of deviations can be generated.

Step 11. Issue Resolution. After all the deviations are generated for all of the base-case scenarios, it is time for the analyst/team to take stock and form some conclusions about their work and the products they've created. They may wish to prioritize the scenarios on a criterion, or quantitatively evaluate the relative likelihood of their occurrence. There may be deviations that are of particular concern, and they may wish to highlight or publicize them for the aviation safety community. It may be decided that some of the scenarios generated need a quantitative analysis to estimate absolute risk. There may be some suggestions or recommendations the team would like to make to reduce the likelihood of certain scenarios and their consequences. This is the place in the prospective analysis to do these kinds of things. Like an executive summary of a long report, the Issue Resolution section should contain the important findings.

CONCLUSIONS

ASHRAM has been developed for anayzing existing accident data and for creating families of accident/incident scenarios based on the premise of the EFC. Using the [detailed versions of the] procedures outlined here, and the forms supplied in [2], aviation safety professionals can systematically generate novel accident/incident scenarios and consider ways of avoiding those scenarios before they become headlines.

The next step for ASHRAM is validation. Sample issues need to be processed so that the technique can be validated and refined. A software implementation of ASHRAM would be advantageous. The author hopes that ASHRAM will become a useful, standard tool in the aviation-safety

professional's toolbox, and that lives will ultimately be saved as a result.

<div align="center">REFERENCES</div>

1. U.S. Nuclear Regulatory Commission, *Technical Basis and Implementation Guidelines for a Technique for Human Event Analysis (ATHEANA)*, NUREG-1624, Rev. 1, U.S. Nuclear Regulatory Commission, September 1999.

2. Miller, D.P. and J.A. Forester, *Aviation Safety Human Reliability Analysis Method (ASHRAM)*, SAND2000-2955, Sandia National Laboratories, Albuquerque, NM, December 2000.

[To order copies of this report, go to orders@ntis.fedworld.gov.]

3. Forsythe, J.C. and C. Wenner, "Surety of human elements of high consequence systems: an organic model," in *Proceedings of the IEA 2000/HFES 2000 Congress*, pp. 3:839-3:842, 2000